

Statistics with Set-Valued Functions

Applications to Inverse Approximate Optimization

Anil Aswani

Abstract Much of statistics relies upon four key elements: a law of large numbers, a calculus to operationalize stochastic convergence, a central limit theorem, and a framework for constructing local approximations. These elements are well-understood for objects in a vector space (e.g., points or functions); however, much statistical theory does not directly translate to sets because they do not form a vector space. Building on probability theory for random sets, this paper uses variational analysis to develop operational tools for statistics with set-valued functions. These tools are first applied to nonparametric estimation (kernel regression of set-valued functions). The second application is to the problem of inverse approximate optimization, in which approximate solutions (corrupted by noise) to an optimization problem are observed and then used to estimate the amount of suboptimality of the solutions and the parameters of the optimization problem that generated the solutions. We show that previous approaches to this problem are statistically inconsistent when the data is corrupted by noise, whereas our approach is consistent under mild conditions.

Keywords set-valued functions · statistics · inverse optimization

1 Introduction

While statistical theory is well-developed for problems concerning (single-valued) functions [13, 42], there has been less work on statistics with sets or set-valued functions. Most attention in statistics on sets has been focused on the problem of estimating a single set under different measurement models [17, 20, 22, 25, 33, 34, 40]. The problem of estimating set-valued functions is less well studied, though it has potential applications in varied domains including healthcare, robotics, and energy. For instance, we study in this paper the problem of inverse approximate optimization, where

This work was supported in part by NSF Award CMMI-1450963.

A. Aswani
Industrial Engineering and Operation Research, University of California, Berkeley, CA, USA
E-mail: aaswani@berkeley.edu

approximate solutions (corrupted by noise) to a parametric optimization problem are observed and then used to estimate the amount of suboptimality of the solutions and the parameters that generated the solutions. Inverse approximate optimization can be used to construct predictive models of human behavior and decision-making, where the explicit model is that an individual makes decisions by approximately solving an optimization problem. Statistical estimation in this context could be used to quantify the tradeoffs made by a particular individual between competing objectives, as well as quantify the predictability of the decision-making process. This particular problem of inverse approximate optimization is related to the broader topic of statistics with set-valued functions because the solution mapping of an (even strictly convex) optimization problem becomes a set when suboptimality of solutions is allowed. Thus a framework for statistics with set-valued functions is needed to study such problems.

A substantial impediment to studying such estimation problems is the lack of statistical tools for random sets and set-valued functions, and two technical issues prevent the use of existing tools. The first is that most statistical theory assumes objects belong to a vector space, which is the case for points and functions. But sets do not form a vector space, and so existing statistical theory cannot be used. This is a fundamental difficulty, and even the usual notion of expectation does not apply to sets [30]. The second is that most statistical theory has been developed by using metrics and distance functions to derive results. But analyzing sets using distances is difficult, and most analysis tools and results for sets do not use this approach [10, 35].

Arguably the most natural approach to statistics with random sets is to define a family of sets parametrized by a random vector, and then perform standard statistical analysis with respect to this parametrization. However it is not clear without further analysis whether stochastic convergence of the estimated parameters implies stochastic convergence of the corresponding set estimates. We study this question in a more general framework and give a counterexample to demonstrate how parameter convergence does not always imply set convergence. Moreover, the parametrization approach does not lead to a useful definition for the expectation of random sets [30]; the reason is that the expectation of the parameters does not characterize the expectation of the set in a way in that ensures the law of large numbers holds.

One goal of this paper is to establish a framework for statistics with set-valued functions, and this requires understanding four main ingredients: a law of large numbers, a calculus to operationalize stochastic convergence, a central limit theorem, and tools for constructing local approximations. Probability theory for random sets [30] provides an expectation for random sets [9, 26], a law of large numbers [3], and a central limit theorem [44]. Here we use variational analysis [35] to advocate a notion of local approximation for set-valued functions, and to develop results that allow us to interpret stochastic convergence and expectations of random sets as operators.

The paper begins by describing our notation and providing some useful definitions related to set-valued functions. We focus in this paper on almost sure (a.s.) convergence because the corresponding definitions and approach most clearly demonstrate the tight link between variational analysis and statistics. Defining set convergence in probability requires metrization, which partially obscures the relationship to variational analysis. We also focus on Lipschitz continuity for set-valued functions because we advocate using this concept as a notion of local approximation for set-

valued functions. The utility of this approach is displayed later in the paper when we use Lipschitz continuity as a replacement for differentiability when proving a Delta method-like result and proving statistical consistency of a kernel regression estimator.

The next section shows how to interpret stochastic convergence and expectation of random sets as operators. We study the limit of sequences of sets under different set operations, after proving a set-based generalization of the continuous mapping theorem [13] from statistics. Then we study the expectation of random sets under various set operations. Standard proofs about the properties of the expectation of random variables do not extend because the expectation of a random set cannot be computed by integration. This means properties like distribution of expectation under independence of the product of a random matrix with a random set or Jensen's inequality have not been previously established, and we prove such results. We conclude by reviewing a law of large numbers and a central limit theorem for random sets.

Another goal of this paper is to study two problems of estimating set-valued functions, and through the process of analyzing these problems we demonstrate the utility of our framework for statistics with set-valued functions. The first problem we study is estimating a set-valued function using noisy measurements of the set. We propose a kernel regression estimator that can be interpreted as a generalization of methods for functions [4, 5, 12, 31, 43]. The key step in proving statistical consistency is using Lipschitz continuity of the set-valued function to construct local approximations. We show that statistical consistency follows by combining our results on stochastic convergence with convergence bounds on (vector-valued) random variables.

The second problem we study is inverse approximate optimization, where noisy measurements of approximate solutions to an optimization problem are used to estimate the amount of suboptimality of the solutions and the parameters of the optimization problem. Inverse optimization is typically studied in the setting of zero measurement noise [1, 15, 19]. Though some methods have been proposed for use with noisy data [11, 23], it was recently shown [8] that these past approaches [11, 23] are statistically inconsistent. This paper studies inverse optimization when the data is simultaneously both noisy and suboptimal, which has not been previously considered in the literature. We construct a new estimator and prove its statistical consistency using the concept of stochastic epi-convergence [4, 18, 21, 24, 27, 38]. Combining this with our results on stochastic convergence and results on the continuity of solutions to optimization problems [35, 36] demonstrates that our approach consistently estimates the (set-valued) approximate solution mapping that generates the data.

We conclude by examining extensions of the problem of inverse approximate optimization, as well as discussing related open questions about statistics with set-valued functions. In particular, we describe how some extensions lead to formulations of optimization problems with structures (e.g., objective functions that are integrals whose domain of integration depends on the decision variable) that have not been well-studied from the perspective of numerical optimization. Performing statistics with sets and set-valued functions also leads to questions about the design of numerical representations of sets. We argue that further study of statistics with set-valued functions will require developing new numerical methods and optimization theory.

2 Preliminaries

This section presents the notation used in this paper, as well as several useful concepts from variational analysis. Most of the variational analysis definitions are from [35]. The definition of set-valued set functions is from [29], and we use the definitions of the Minkowski set operations from [39]. The as-lim inf, as-lim sup notations are new.

2.1 Notation

Let $\mathcal{F}(E)$ be the space of closed subsets of E , and let $\mathcal{K}(E)$ be the space of compact subsets of E . We will focus on cases where E is a Euclidean space, and so will use the notation \mathcal{F}, \mathcal{K} to refer to the corresponding spaces. Clearly $\mathcal{F} \supset \mathcal{K}$ by definition.

Suppose C, D are sets and Ψ is a matrix or scalar. Then we use the following set notation: $C \cup D$ is the union of C, D ; $C \cap D$ is the intersection of C, D ; $C \subseteq D$ denotes that C is a subset of D ; $C \supseteq D$ denotes that C is a superset of D ; $\text{cl}(C)$ is the closure of C ; $\text{co}(C)$ is the convex hull of C ; C^c is the complement of C ; ∂C is the boundary of C ; $C \oplus D = \{c + d : c \in C, d \in D\}$ is the Minkowski sum of C and D ; $C \ominus D = \{x : x \oplus D \subseteq C\}$ is the Minkowski difference of C and D ; and $\Psi \cdot C = \{\Psi \cdot c : c \in C\}$.

2.2 Limit Definitions and Stochastic Convergence

The outer limit of the sequence of sets C^ν is defined as

$$\limsup_\nu C^\nu = \{x : \exists \nu_k \text{ s.t. } x^{\nu_k} \rightarrow x \text{ with } x^{\nu_k} \in C^{\nu_k}\}, \quad (1)$$

and the inner limit of the sequence of sets C^ν is defined as

$$\liminf_\nu C^\nu = \{x : \exists x^\nu \rightarrow x \text{ with } x^\nu \in C^\nu\}. \quad (2)$$

The outer limit consists of all the cluster points of C_n , whereas the inner limit consists of all limit points of C_n . The limit of the sequence of sets C^ν exists if the outer and inner limits are equal, and we define that $\lim_\nu C^\nu := \limsup_\nu C^\nu = \liminf_\nu C^\nu$.

Let $\overline{\mathbb{R}} = [-\infty, \infty]$ denote the extended real line. A sequence of extended-real-valued functions $f^\nu : X \rightarrow \overline{\mathbb{R}}$ is said to epi-converge to f if at each $x \in X$ we have

$$\begin{cases} \liminf_\nu f^\nu(x^\nu) \geq f(x) & \text{for every sequence } x^\nu \rightarrow x \\ \limsup_\nu f^\nu(x^\nu) \leq f(x) & \text{for some sequence } x^\nu \rightarrow x \end{cases} \quad (3)$$

The notion of epi-convergence is so-named because it is equivalent to set convergence of the epigraphs of f^ν , meaning that epi-convergence is equivalent to the condition $\lim_\nu \{(x, \alpha) \in X \times \mathbb{R} : f^\nu(x) \leq \alpha\} = \{(x, \alpha) \in X \times \mathbb{R} : f(x) \leq \alpha\}$.

We also define compact notation to specify stochastic convergence. The notation $\text{as-lim sup}_\nu C^\nu \subseteq C$ denotes $\mathbb{P}(\limsup_\nu C^\nu \subseteq C) = 1$, the notation $\text{as-lim inf}_\nu C^\nu \supseteq C$ denotes $\mathbb{P}(\liminf_\nu C^\nu \supseteq C) = 1$, and the notation $\text{as-lim}_\nu C^\nu = C$ denotes $\mathbb{P}(\lim_\nu C^\nu = C) = 1$. Note $\text{as-lim sup}_\nu C^\nu \subseteq C$ and $\text{as-lim inf}_\nu C^\nu \supseteq C$ if and only if $\text{as-lim}_\nu C^\nu = C$, since a countable intersection of almost sure events occurs almost surely.

2.3 Set-Valued Mappings

A set-valued set function $G : V \rightrightarrows U$ assigns to each set $S \subseteq V$ a set $G(S) \subseteq U$. The outer limit of G at the set $\bar{S} \in V$ is defined as

$$\limsup_{S \rightarrow \bar{S}} G(S) = \{u : \exists S^V \rightarrow \bar{S} \text{ s.t. } u^V \rightarrow u \text{ with } S^V \subseteq V, u^V \in G(S^V)\}, \quad (4)$$

and the inner limit of G at the set $\bar{S} \subseteq V$ is defined as

$$\liminf_{S \rightarrow \bar{S}} G(S) = \{u : \forall S^V \rightarrow \bar{S}, \exists u^V \rightarrow u \text{ with } S^V \subseteq V, u^V \in G(S^V)\}. \quad (5)$$

The intuition is similar to the notions for sequences of sets. The set-valued set function G is outer semicontinuous (osc) at \bar{S} if $\limsup_{S \rightarrow \bar{S}} G(S) \subseteq G(\bar{S})$, and G is inner semicontinuous (isc) at \bar{S} if $\liminf_{S \rightarrow \bar{S}} G(S) \supseteq G(\bar{S})$. The set-valued set function G is continuous at \bar{S} when it is both osc and isc, that is when $\lim_{S \rightarrow \bar{S}} G(S) = G(\bar{S})$.

Variational analysis typically uses set-valued functions, rather than set-valued set functions. A set-valued function $F : X \rightrightarrows U$ assigns to each point $x \in X$ a set $F(x) \subseteq U$. Outer limits, inner limits, outer semicontinuity, inner semicontinuity, and continuity are defined as above but with points replacing sets in the domain. Moreover, a set-valued function applied pointwise to sets is an osc, isc, continuous set-valued set function whenever the set-valued function is osc, isc, continuous, respectively.

2.4 Distances and Lipschitz Continuity

Let $d(x, C) = \inf_{y \in C} \|x - y\|$ and $\rho(x, C) = \inf_{y \in C} \|x - y\|^2$ be the distance and squared distance, respectively, from a point x to set C . The support function of C is $h(x, C) = \sup_{y \in C} x^\top y$. We also define the indicator function $\delta(x, C)$ to equal 0 when $x \in C$ and $+\infty$ when $x \notin C$. The (integrated) set distance between C and D is defined as

$$d_l(C, D) = \int_0^\infty d_r(C, D) e^{-r} dr, \quad (6)$$

where the psuedo-distance between sets C and D is given by

$$d_r(C, D) = \max_{\|x\| \leq r} |d(x, C) - d(x, D)|. \quad (7)$$

The integrated set distance d_l is a metric that characterizes the convergence defined earlier for sets in \mathcal{F} , and the Pompeiu-Hausdorff distance d_∞ is a metric that characterizes the convergence defined earlier for sets in \mathcal{K} . Since these metrics are complex, the sequence characterization of convergence is arguably more natural for sets.

One exception to this statement is in defining Lipschitz continuity for set-valued functions. A set-valued function $F : X \rightrightarrows U$ is Lipschitz continuous on X with constant $\kappa \in \mathbb{R}_+$ if it is nonempty-closed-valued and such that

$$F(x') \subseteq F(x) + \kappa \|x' - x\| \mathbb{B} \text{ for } x, x' \in X, \quad (8)$$

where $\mathbb{B} = \{u : \|u\| \leq 1\}$ is the unit ball. A set-valued set function $G : V \Rightarrow U$ is Lipschitz continuous on V with constant $\kappa \in \mathbb{R}_+$ if it is nonempty-closed-valued and

$$G(S') \subseteq G(S) + \kappa d_\infty(S', S) \mathbb{B} \text{ for } S, S' \in V. \quad (9)$$

We will make use of Lipschitz continuity as a zeroth-order local approximation.

3 Mathematical Tools for Statistics with Set-Valued Functions

This section develops mathematical tools that allow us to interpret stochastic convergence and the expectation of random sets as operators. We prove results on the limit of sequences of sets under different set operations, define an expectation for random sets, and then derive results about the behavior of this expectation under different set operations. We conclude this section by briefly summarizing a large of large numbers and a central limit theorem for random sets.

3.1 Stochastic Limit Theorems

Our reason for considering set-valued set functions is this allows us to more precisely generalize the classical *continuous mapping theorem* of statistics [13] to mappings applied to sequences of sets. Because semicontinuity is an important aspect of set convergence, a generalization that considers semicontinuity leads to a richer set of results than simply considering continuity.

Theorem 1 (Semicontinuous Mapping Theorem) *Let G be a set-valued set function, and suppose $\text{as-lim}_n C_n = C$. There are three cases:*

- (a) *If G is osc at C , then $\text{as-lim sup}_n G(C_n) \subseteq G(C)$.*
- (b) *If G is isc at C , then $\text{as-lim inf}_n G(C_n) \supseteq G(C)$.*
- (c) *If G is continuous at C , then $\text{as-lim}_n G(C_n) = G(C)$.*

Proof The definition of osc (isc) means $\lim_n C_n = C$ implies $\limsup_n G(C_n) \subseteq G(C)$ ($\liminf_n G(C_n) \supseteq G(C)$). This means $\mathbb{P}(\limsup_n G(C_n) \subseteq G(C)) \geq \mathbb{P}(\lim_n C_n = C) = 1$ ($\mathbb{P}(\liminf_n G(C_n) \supseteq G(C)) \geq \mathbb{P}(\lim_n C_n = C) = 1$), which shows the first two cases. The third case follows from the first two cases by recalling that continuity at C is equivalent to being both osc and isc at C . \square

Remark 1 One consequence is that a set $S(\theta)$ parametrized by θ has the behavior that $\text{as-lim}_n \theta_n = \theta_0$ implies $\text{as-lim}_n S(\theta_n) = S(\theta_0)$ only when the set is continuous with respect to the parametrization. For example, consider the parametrized set $S(\theta) = \{1\}$ if $\theta > 0$, $S(\theta) = \{-1\}$ if $\theta < 0$, and $S(\theta) = [-1, 1]$ if $\theta = 0$. If $\theta_n = 1/n$, then $S(\theta_n) \equiv \{1\}$ and so $\text{as-lim}_n S(\theta_n) = \{1\}$. But $\text{as-lim}_n \theta_n = 0$ and $S(0) = [-1, 1]$.

As is customary in statistics, we immediately get some useful corollaries to our semicontinuous mapping theorem by applying the theorem to specific mappings. Our first corollary applies the semicontinuous mapping theorem to set operations like unions and intersections of sets, the boundary of sets, the convex hull of sets, etc.

Corollary 1 Let $C_n, D_n \in \mathcal{F}$ be almost surely convergent sequences of sets (i.e., $\text{as-lim}_n C_n = C$ and $\text{as-lim}_n D_n = D$). Then we have:

- (a) $\text{as-lim}_n (C_n \cup D_n) = C \cup D$
- (b) $\text{as-lim sup}_n (C_n \cap D_n) \subseteq C \cap D$
- (c) $\text{as-lim inf}_n \text{cl}(C_n^c) \supseteq \text{cl}(C^c)$
- (d) $\text{as-lim inf}_n \partial C_n \supseteq \partial C$
- (e) $\text{as-lim inf}_n \text{co}(C_n) \supseteq \text{co}(C)$
- (f) $\text{as-lim}_n \text{co}(D_n) = \text{co}(D)$

Proof We interpret \cup, \cap as set-valued set functions with a domain over the product space $\mathcal{F} \times \mathcal{F}$: The function $G_1(S, T) = S \cup T$ is continuous [29], and the function $G_2(S, T) = S \cap T$ is osc [29]. The set complement and boundary operators can be interpreted as set-valued set functions with domain \mathcal{F} : The function $G_3(S) = \text{cl}(S^c)$ is isc [29], and the function $G_4(S) = \partial S$ is isc [29]. The convex hull operation can be cast as set-valued set functions: $G_5(S) = \text{co}(S)$ is isc when the domain is \mathcal{F} , and $G_6(S) = \text{co}(S)$ is continuous when the domain is \mathcal{K} [29]. The results now follow from the corresponding parts of the semicontinuous mapping theorem. \square

Remark 2 Note the above result states that the stochastic limit of the convex hull operator is sensitive to the domain of the sequence of sets.

We can also apply the semicontinuous mapping theorem to the Minkowski set operations. These results are useful for proving convergence of statistical estimators.

Corollary 2 Let $C_n \in \mathcal{F}$ and $D_n \in \mathcal{K}$ be almost surely convergent sequences of sets (i.e., $\text{as-lim}_n C_n = C$ and $\text{as-lim}_n D_n = D$), and let Ψ_n be an almost surely convergent sequence of matrices or scalars (i.e., $\text{as-lim}_n \Psi_n = \Psi$). Then

- (a) $\text{as-lim}_n (C_n \oplus D_n) = C \oplus D$
- (b) $\text{as-lim sup}_n (C_n \ominus D_n) \subseteq C \ominus D$, when $D \neq \emptyset$
- (c) $\text{as-lim}_n \Psi_n \cdot C_n = \Psi \cdot C$
- (d) $\text{as-lim}_n \Psi_n^{-1} C_n = \Psi^{-1} C$, when Ψ is invertible

Proof We interpret \oplus, \ominus as set-valued set functions with a domain over the product space $\mathcal{F} \times \mathcal{K}$: The function $G_1(S, T) = S \oplus T$ is continuous [29], and the function $G_2(S, T) = S \ominus T$ is osc if $T \neq \emptyset$ [29]. So the first two results follow from Theorem 1. The multiplication operation can be interpreted as a set-valued set function $G_3(S, T) = T \cdot S$ with domain over the product space $\mathbb{M} \times \mathcal{F}$, where \mathbb{M} is the space of matrices of appropriate dimension or the space of scalars. We show it is continuous. Suppose G_3 is not osc at $\bar{S} \times \bar{T}$; then there exist $T^\nu \rightarrow \bar{T}$, $S^\nu \rightarrow \bar{S}$, and $u^\nu \rightarrow \bar{u}$ with $T^\nu \in \mathbb{M}$, $S^\nu \in \mathcal{F}$, $u^\nu \in T^\nu \cdot S^\nu$, and $\bar{u} \notin \bar{T} \cdot \bar{S}$. But by the definition of matrix-set (or scalar-set) multiplication there exists $v^\nu \in S^\nu$ with $u^\nu = T^\nu \cdot v^\nu$, and by the outer limit definition for a sequence of sets there exist v_k and \bar{v} such that $v^{v_k} \rightarrow \bar{v}$ with $\bar{v} \in \bar{S}$, which is a contradiction since matrix-vector (or scalar-vector) multiplication is osc. Thus G_3 is osc. Next, we show G_3 is isc at $\bar{T} \cdot \bar{S}$: Consider any $\bar{x} \in \bar{S}$ and $u = \bar{T} \cdot \bar{x}$, and let T^ν, S^ν be any sequences satisfying $T^\nu \rightarrow \bar{T}$ and $S^\nu \rightarrow \bar{S}$. By the inner limit definition there exists $x^\nu \rightarrow \bar{x}$ with $x^\nu \in S^\nu$, and so $T^\nu \cdot x^\nu \rightarrow \bar{T} \cdot \bar{x}$ with $T^\nu \cdot x^\nu \in T^\nu \cdot S^\nu$. So G_3 satisfies the definition of being isc at $\bar{T} \cdot \bar{S}$, and is continuous since it is also osc.

The third result follows from Theorem 1. The fourth result is proved by noting Theorem 1 implies $\text{as-lim}_n \Psi_n^{-1} = \Psi^{-1}$ since the matrix inverse operation is continuous except at points of singularity, and so $\text{as-lim}_n \Psi_n^{-1} C_n = \Psi^{-1} C$ by the third result. \square

Our final results on stochastic limits are not based on the semicontinuous mapping theorem, but are nevertheless useful for writing stochastic convergence proofs.

Lemma 1 (Sandwich Lemma) *Let $L_n \in \mathcal{F}$ and $U_n \in \mathcal{F}$ be almost surely convergent sequences of sets (i.e., $\text{as-lim}_n L_n = L$ and $\text{as-lim}_n U_n = U$), and let $C_n \in \mathcal{F}$ be a sequence of sets. Then we have*

- (a) $\text{as-lim sup}_n C_n \subseteq U$, when $C_n \subseteq U_n$
- (b) $\text{as-lim inf}_n C_n \supseteq L$, when $C_n \supseteq L_n$
- (c) $\text{as-lim}_n C_n = L = U$, when $L_n \subseteq C_n \subseteq U_n$ and $L = U$

Proof For the first two results, note $\text{as-lim sup}_n C_n \subseteq \text{as-lim sup}_n U_n = \text{as-lim}_n U_n = U$ and $\text{as-lim inf}_n C_n \supseteq \text{as-lim inf}_n L_n = \text{as-lim}_n L_n = L$. The third result follows from the first two results and the definition of limit. \square

This sandwich lemma is not only valuable for statistical analysis, but also leads to a convergence result that is helpful in proving statistical consistency.

Corollary 3 *Let $C_n \in \mathcal{F}$ and $D_n \in \mathcal{K}$ be sequences of sets. If $\text{as-lim}_n D_n = \{0\}$ and $\text{as-lim}_n C_n \oplus D_n$ exists, then $\text{as-lim}_n C_n = \text{as-lim}_n C_n \oplus D_n$.*

Proof Since $C_n \subseteq C_n \oplus D_n$, the sandwich lemma implies $\text{as-lim sup}_n C_n \subseteq \text{as-lim}_n C_n \oplus D_n$. Let $u_n \in C_n \oplus D_n$ be any almost surely convergent sequence. Then by definition there exist $c_n \in C_n$ and $d_n \in D_n$ such that $u_n = c_n + d_n$. So $|u_n - c_n| = |d_n|$, and hence $\text{as-lim}_n u_n = \text{as-lim}_n c_n$ since $\text{as-lim}_n D_n = \{0\}$ by assumption. Thus $\text{as-lim inf}_n C_n \supseteq \text{as-lim}_n C_n \oplus D_n$, and the result follows by again applying the sandwich lemma. \square

Remark 3 Note that $\text{as-lim}_n C_n$ does not necessarily exist if $\text{as-lim}_n C_n \oplus D_n$ exists but $\text{as-lim}_n D_n \neq \{0\}$. For instance, consider the deterministic sequences $C_n = \{(-1)^n\}$ and $D_n = \{(-1)^{n+1}\}$. Then $C_n \oplus D_n \equiv \{0\}$, but C_n and D_n clearly do not converge.

3.2 Expectation

Because sets do not form a vector space, defining expectations for random sets is not straightforward. In fact, a number of different definitions have been proposed [30] that capture different features that might be desired for an expectation operation. One particularly useful definition is the *selection expectation* [26, 9]. This definition for the expectation of random sets is the most well studied because it leads to a corresponding law of large numbers and central limit theorem [30].

For a random set X , a selection ξ is a (single-valued) random vector that almost surely belongs to X . We say the selection ξ is integrable if $\mathbb{E}\|\xi\|_1$ is finite, where $\|\cdot\|_1$ is the usual ℓ_1 -norm. The selection expectation of a random set X is defined as

$$\mathbb{E}(X) = \text{cl}\{\mathbb{E}\xi : \xi \in \mathcal{S}^1(X)\}, \quad (10)$$

where $\mathcal{S}^1(X)$ is the set of all integrable selections of X . The random set X is called integrable if $\mathcal{S}^1(X) \neq \emptyset$, and note this property implies X is almost surely non-empty.

The selection expectation is difficult to use because it cannot be computed by taking an integral, as is the case for expectations for objects in a vector space. But since we assume E is Euclidean space, the definition of the selection expectation simplifies and has a sharp characterization [30]: If the probability space is nonatomic and X is a bounded and closed integrable random set, then $\mathbb{E}(X) = \{\mathbb{E}\xi : \xi \in \mathcal{S}^1(X)\}$ is a compact set, $\mathbb{E}(X)$ is convex, $\mathbb{E}(X) = \mathbb{E}(\text{co}(X))$, and $h(u, \mathbb{E}(X)) = \mathbb{E}(h(u, X))$ for all $u \in E$. This support function characterization is powerful, and allows us to prove a number of helpful properties about the selection expectation. More importantly, the following results allow us to operationalize the selection expectation, which is useful from a practical standpoint for performing statistical analysis.

Proposition 1 *Suppose C, D are bounded and closed integrable random sets, and let Ψ be a random matrix or a random scalar. If the probability space is nonatomic, then*

- (a) $\mathbb{E}(C) = \text{co}(C)$, when C is deterministic
- (b) $\mathbb{E}(C \oplus D) = \mathbb{E}(C) \oplus \mathbb{E}(D)$
- (c) $\mathbb{E}(\Psi C) = \mathbb{E}(\Psi) \cdot \mathbb{E}(C)$, when Ψ is independent of C
- (d) $\mathbb{E}(C) \subseteq \mathbb{E}(D)$, when $C \subseteq D$
- (e) $\mathbb{E}(C) \cup \mathbb{E}(D) \subseteq \mathbb{E}(C \cup D)$
- (f) $\mathbb{E}(C \cap D) \subseteq \mathbb{E}(C) \cap \mathbb{E}(D)$
- (g) $\mathbb{E}(C \ominus D) \subseteq \mathbb{E}(C) \ominus \mathbb{E}(D)$, when $C \ominus D$ is non-empty.

Proof The first result holds since $\mathbb{E}(X) = \mathbb{E}(\text{co}(X))$ and $h(u, \mathbb{E}(C)) = \mathbb{E}(h(u, C)) = h(u, C)$. The next result follows from $h(u, C \oplus D) = h(u, C) + h(u, D)$ [39], since $h(u, \mathbb{E}(C \oplus D)) = \mathbb{E}(h(u, C \oplus D)) = \mathbb{E}(h(u, C) + h(u, D)) = \mathbb{E}(h(u, C)) + \mathbb{E}(h(u, D)) = h(u, \mathbb{E}(C)) + h(u, \mathbb{E}(D)) = h(u, \mathbb{E}(C) \oplus \mathbb{E}(D))$. The fourth result holds since $h(u, C) \leq h(u, D)$ when $C \subseteq D$ [39], which implies $h(u, \mathbb{E}(C)) = \mathbb{E}(h(u, C)) \leq \mathbb{E}(h(u, D)) = h(u, \mathbb{E}(D))$. For the fifth result, note $C \subseteq C \cup D$ and $D \subseteq C \cup D$. The fourth result gives $\mathbb{E}(C) \subseteq \mathbb{E}(C \cup D)$ and $\mathbb{E}(D) \subseteq \mathbb{E}(C \cup D)$, which implies $\mathbb{E}(C) \cup \mathbb{E}(D) \subseteq \mathbb{E}(C \cup D)$. The sixth result follows since combining $C \cap D \subseteq C$, $C \cap D \subseteq D$, and the fourth result gives: $\mathbb{E}(C \cap D) \subseteq \mathbb{E}(C)$ and $\mathbb{E}(C \cap D) \subseteq \mathbb{E}(D)$, which implies $\mathbb{E}(C \cap D) \subseteq \mathbb{E}(C) \cap \mathbb{E}(D)$. To prove the seventh result, note $(C \ominus D) \oplus D \subseteq C$ [39]. Applying the second and fourth results yields $\mathbb{E}(C \ominus D) \oplus \mathbb{E}(D) \subseteq \mathbb{E}(C)$, and so $\mathbb{E}(C \ominus D) \subseteq \mathbb{E}(C) \ominus \mathbb{E}(D)$ [39].

The third result cannot be proved using support functions since Ψ can have negative terms. Our approach is to show $\mathcal{S}^1(\Psi C) = \Psi \mathcal{S}^1(C)$, since we would then have $\mathbb{E}(\Psi C) = \{\mathbb{E}(\Psi \xi) : \xi \in \mathcal{S}^1(C)\} = \{\mathbb{E}(\Psi) \cdot \mathbb{E}(\xi) : \xi \in \mathcal{S}^1(C)\} = \mathbb{E}(\Psi) \cdot \{\mathbb{E}(\xi) : \xi \in \mathcal{S}^1(C)\} = \mathbb{E}(\Psi) \cdot \mathbb{E}(C)$. The inclusion $\mathcal{S}^1(\Psi C) \supseteq \Psi \mathcal{S}^1(C)$ is obvious by definition. To prove the reverse inclusion, let $\{\xi_n, n \geq 1\}$ with $\xi_n \in \mathcal{S}^1(C)$ be the Castaing representation [30] of C . Then $\{\Psi \xi_n, n \geq 1\}$ is the Castaing representation of ΨC . But by Lemma 1.3 of [30], each selection in $\mathcal{S}^1(\Psi C)$ can be approximated arbitrarily well by step functions with arguments from $\{\Psi \xi_n, n \geq 1\}$. This means $\mathcal{S}^1(\Psi C) \subseteq \Psi \mathcal{S}^1(C)$, and so $\mathcal{S}^1(\Psi C) = \Psi \mathcal{S}^1(C)$ since both inclusions were shown. \square

Remark 4 Note the assumptions for part (b) include the cases where: Ψ is deterministic, C is deterministic, or Ψ has positive or negative entries.

Another result used in statistics is Jensen's inequality [13], which bounds changing the order of applying an expectation and a convex function to a random variable. Our next result shows we can generalize Jensen's inequality to set-valued functions.

Proposition 2 (Jensen's Inequality) *Let $S(u)$ be a graph-convex set-valued function (i.e., $S((1-\lambda)u_0 + \lambda u_1) \supseteq (1-\lambda) \cdot S(u_0) + \lambda \cdot S(u_1)$ for $\lambda \in (0,1)$), and let X be bounded and closed integrable random set. If $S(\cdot)$ is locally bounded (i.e., $S(B)$ is bounded for every bounded set B) and continuous, then we have $S(\mathbb{E}(X)) \supseteq \mathbb{E}(S(X))$.*

Proof The selection expectation equals the Debreu expectation under our assumptions [30]. This means there exists a sequence of random sets X_n with the distribution

$$X_n = F_{in} \text{ with probability } p_{in}, \text{ for } i = 1, \dots, n \quad (11)$$

such that $\text{as-lim } X_n = X$, $\mathbb{E}(X) = \lim_n \mathbb{E}(X_n)$, and $\mathbb{E}(X_n) = \bigoplus_{i=1}^n p_{in} \cdot F_{in}$. Using the semicontinuous mapping theorem implies $\text{as-lim}_n S(X_n) = S(X)$, and so we have equality of the selection expectation and Debreu expectation [30]. This means that $\mathbb{E}(S(X)) = \lim_n \mathbb{E}(S(X_n))$ and $\mathbb{E}(S(X_n)) = \bigoplus_{i=1}^n p_{in} \cdot S(F_{in})$. Next note $S(\bigoplus_{i=1}^n p_{in} \cdot F_{in}) \supseteq \bigoplus_{i=1}^n p_{in} \cdot S(F_{in})$ by the graph-convexity of $S(\cdot)$. Taking the limit of this set relationship gives $S(\mathbb{E}(X)) = \lim S(\bigoplus_{i=1}^n p_{in} \cdot F_{in}) \subseteq \lim_n \bigoplus_{i=1}^n p_{in} \cdot S(F_{in}) = \mathbb{E}(S(X))$, where we have used the fact that $\lim_n S(\bigoplus_{i=1}^n p_{in} \cdot F_{in}) = S(\mathbb{E}(X))$ by definition of the continuity of the set-valued function $S(\cdot)$. \square

Remark 5 Jensen's inequality is sometimes stated for concave functions, and such a generalization exists for set-valued mappings. If $S(u)$ is a graph-concave set-valued function (i.e., $S((1-\lambda)u_0 + \lambda u_1) \subseteq (1-\lambda) \cdot S(u_0) + \lambda \cdot S(u_1)$ for $\lambda \in (0,1)$) and the other assumptions of the above theorem hold, then we have $S(\mathbb{E}(X)) \subseteq \mathbb{E}(S(X))$.

Lastly, we present a strong law of large numbers (SLLN) for the selection expectation. The key idea is the Minkowski sum takes the role of averaging.

Theorem 2 (Artstein and Vitale, 1975 [3]) *Suppose the probability space is nonatomic. If X, X_i , $i \geq 1$, are i.i.d. bounded and closed integrable random sets, then we have that: $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n X_i = \mathbb{E}(X)$.*

This particular strong law of large numbers can be generalized in a number of ways, and a survey of the different generalizations possible can be found in [30].

3.3 Central Limit Theorems

Unlike laws of large numbers that relate convergence of Minkowski sums of i.i.d. random variables $\frac{1}{n} \bigoplus_{i=1}^n X_i$ to their selection expectation $\mathbb{E}(X)$, analogs of the central limit theorem (CLT) relating Minkowski sums and selection expectations are less well-understood. One major impediment is that the \ominus operator does not generally invert the \oplus operator, which means it is generally not possible to normalize (in the sense of having a zero mean) the Minkowski sum $\frac{1}{n} \bigoplus_{i=1}^n X_i$. As a result, the standard approach to generalizing the central limit theorem is to normalize by instead considering the Hausdorff distance between Minkowski sum and the selection expectation.

Theorem 3 (Weil, 1982 [44]) Suppose the probability space is nonatomic. If $X, X_i, i \geq 1$, are i.i.d. bounded and closed integrable random sets, then we have that: $\sqrt{n} \cdot d_\infty(\frac{1}{n} \bigoplus_{i=1}^n X_i, \mathbb{E}(X)) \rightarrow \sup_u \{\|\zeta(u)\| \mid \|u\| \leq 1\}$ in distribution, where $\zeta(u)$ for $\|u\| \leq 1$ is a centered Gaussian random field with covariance given by: $\mathbb{E}(\zeta(u) \cdot \zeta(v)) = \mathbb{E}(h(u, X) \cdot h(v, X)) - \mathbb{E}(h(u, X)) \cdot \mathbb{E}(h(v, X))$.

The difficulty with this central limit theorem is that it lacks a clear geometrical interpretation (in contrast to the classical central limit theorem for random variables) for the limiting distribution, and the question of whether such a geometrical interpretation exists remains an open question [30]. However, one advantage of this formulation is that it lends itself to a generalization of the Delta method [13] from statistics.

Proposition 3 (Approximate Delta Method) Suppose $r_n d_\infty(C_n, C) \rightarrow w$ in distribution, where r_n is a strictly increasing sequence, C_n is a sequence of random sets, C is a deterministic set, and w is a random variable. If S is a Lipschitz continuous set-valued set function, then

$$\limsup_n \mathbb{P}(r_n d_\infty(S(C_n), S(C)) \geq u) \leq \mathbb{P}(\kappa \cdot w \geq u), \quad (12)$$

where $\kappa \in \mathbb{R}_+$ is the Lipschitz constant of S .

Proof Lipschitz continuity of S gives $r_n d_\infty(S(C_n), S(C)) \leq \kappa \cdot r_n d_\infty(C_n, C)$. Thus $\mathbb{P}(r_n d_\infty(S(C_n), S(C)) \geq u) \leq \mathbb{P}(\kappa \cdot r_n d_\infty(C_n, C) \geq u)$. The limit superior of both sides gives the result since $r_n d_\infty(C_n, C) \rightarrow w$ in distribution. \square

Remark 6 The Delta method relates asymptotic distributions of random variables under differentiable functions [13], and the intuition is the derivative is used as a local approximation of the function. The above result demonstrates one instance where Lipschitz continuity can be used as a local approximation for set-valued mappings.

Though \ominus does not generally invert \oplus , there is one special case when inversion is possible. If C, D are compact convex sets, then $(C \oplus D) \ominus D = C$ [39]. Using this property, we describe a new central limit theorem for random sets with a particular structure that is useful for statistical applications. Specifically, this result applies to randomly translated sets (RaTS), which are random sets of the form $C = K \oplus \xi$, where K is a deterministic compact convex set, and ξ is a (vector-valued) random variable.

Theorem 4 (Central Limit Theorem for RaTS) Suppose the probability space is nonatomic, and that $X, X_i, i \geq 1$, are i.i.d. random sets with $X_i = K \oplus \xi_i$, where K is a deterministic compact convex set and $\xi, \xi_i, i \geq 1$, are i.i.d. (vector-valued) random variables with zero mean and finite variance. Then

$$\sqrt{n} \cdot ((\frac{1}{n} \bigoplus_{i=1}^n X_i) \ominus \mathbb{E}(X)) \rightarrow \mathcal{N}(0, \mathbb{E}(\xi \xi^\top)) \quad (13)$$

in distribution, where $\mathcal{N}(0, \mathbb{E}(\xi \xi^\top))$ is a jointly Gaussian random variable with zero mean and covariance matrix given by $\mathbb{E}(\xi \xi^\top)$.

Proof Since $(\frac{1}{n} \bigoplus_{i=1}^n X_i) \ominus \mathbb{E}(X) = ((\frac{1}{n} \sum_{i=1}^n \xi_i) \oplus K) \ominus K = \frac{1}{n} \sum_{i=1}^n \xi_i$, the result follows by the classical central limit theorem [13]. \square

The benefit of this new formulation of the central limit theorem is that it has a clear geometrical interpretation like the classical central limit theorem for random variables, but unfortunately this result only applies to the specific class of RaTS.

4 Kernel Regression

We will construct a nonparametric estimator for set-valued functions using an approach that can be viewed as a natural generalization of kernel regression methods for functions [4, 5, 12, 31, 43]. These techniques are considered nonparametric because no assumptions (beyond minimal assumptions about regularity) are made about the functional relationship between the input and outputs of the function being estimated.

4.1 Problem Setup

Consider a Lipschitz continuous set-valued function $S(u) : X \rightrightarrows \mathbb{R}^q$ with data $(x_i, S_i) \in X \times \mathbb{R}^q$ for $i = 1, \dots, n$, where: $X \subseteq \mathbb{R}^d$ is a compact set; $S(u)$ is a convex compact set for each $u \in X$; x_i are i.i.d. (vector-valued) random variables with a Lipschitz continuous density function f_X that has the property $f_X(u) > 0$ for $u \in X$; and $S_i = S(x_i) \oplus w_i$ with w_i are i.i.d. (vector-valued) random variables that have zero mean $\mathbb{E}(w) = 0$ and finite variance $\|\mathbb{E}(ww^\top)\| < +\infty$.

4.2 Kernel Functions

Kernel regression is so named because these approaches use kernel functions $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, which are functions that are non-negative, bounded, even (i.e., $\varphi(-u) = \varphi(u)$), and have finite support (i.e., there is a constant $\eta \in (0, 1)$ such that $\varphi(u) > 0$ when $|u| \leq \eta$, and $\varphi(u) = 0$ for $|u| \geq 1$). One example of a kernel function is the indicator function $\varphi(u) = (1/2) \cdot \mathbf{1}(|u| < 1)$, and another example is the Epanechnikov kernel $\varphi(u) = (3/4) \cdot (1 - u^2) \cdot \mathbf{1}(u \leq 1)$. Notationally, it is useful to define the family of kernel functions $\varphi_h(u) = h^{-d} \varphi(\|u\|/h)$ and the constant $\gamma = \int_{\mathbb{R}^d} \varphi_1(z) dz$. We first prove a lemma concerning this family of kernel functions.

Lemma 2 *If $h = n^{-1/(d+4)}$, then for $u \in \text{int}(X)$ we have*

- (a) $\text{as-lim}_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) = \gamma \cdot f_X(u)$
- (b) $\text{as-lim}_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot w_i = 0$
- (c) $\text{as-lim}_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \|x_i - u\| = 0$

Proof To prove the first result, observe that

$$\lim_n \sum_{i=1}^n n^{-2} \cdot \text{var}(\varphi_h(x_i - u)) \leq \lim_n c / (nh^d) < \infty, \quad (14)$$

where the first inequality holds for some constant $c \in \mathbb{R}^+$ because φ_h is bounded, non-negative, and has a finite support; and the second inequality holds because $nh^d = n^{4/(d+4)}$. The finiteness of the above summation means we can apply Kolmogorov's strong law of large numbers, which gives $\text{as-lim}_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) - \mathbb{E}(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u)) = 0$. Our next step is to compute this expectation. Note that

$$\begin{aligned} \mathbb{E}(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u)) &= \int_{X \in \mathbb{R}^d} \varphi_h(x - u) \cdot f_X(x) dx \\ &= \int_{z \in \mathbb{R}^d} \varphi_1(z) \cdot f_X(u + hz) dz \end{aligned} \quad (15)$$

where in the last line we made the change of variables $z = (x - u)/h$. So

$$\begin{aligned}
& \left| \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \right) - f_X(u) \cdot \int_{z \in \mathbb{R}^d} \varphi_1(z) dz \right| \\
& \leq \int_{z \in \mathbb{R}^d} \varphi_1(z) \cdot |f_X(u + hz) - f_X(u)| dz \\
& \leq \int_{z \in \mathbb{R}^d} \varphi_1(z) \cdot \kappa h \|z\| dz \\
& \leq h \cdot \kappa \int_{z \in \mathbb{R}^d} \varphi_1(z) dz
\end{aligned} \tag{16}$$

where $\kappa \in \mathbb{R}_+$ is the Lipschitz constant of the density $f_X(u)$. The finite support of the kernel implies the third inequality and that $\int_{z \in \mathbb{R}^d} \varphi_1(z) dz$ is finite. Thus taking the limit of the above inequalities gives $\lim_n \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \right) = f_X(u) \cdot \int_{z \in \mathbb{R}^d} \varphi_1(z) dz$. This proves the first result when combined with the implication of Kolmogorov's strong law of large numbers in our setting.

For the proof of the second result, let $\langle w \rangle_j$ denote the j -th component of the vector w . Next observe that

$$\lim_n \sum_{i=1}^n n^{-2} \cdot \text{var}(\varphi_h(x_i - u) \cdot \langle w_i \rangle_j) \leq \lim_n c \cdot \text{var}(\langle w \rangle_j) / (nh^d) < \infty, \tag{17}$$

where the first inequality holds for some constant $c \in \mathbb{R}^+$ because the w_i have zero mean and because φ_h is bounded, non-negative, and has a finite support; and the second inequality holds because $nh^d = n^{4/(d+4)}$. The finiteness of the above summation means Kolmogorov's strong law of large numbers gives $\text{as-}\lim_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \langle w_i \rangle_j - \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \langle w_i \rangle_j \right) = 0$. But the w_i are zero mean, and so we have that $\mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \langle w_i \rangle_j \right) = 0$.

To prove the third result, observe that

$$\lim_n \sum_{i=1}^n n^{-2} \cdot \text{var}(\varphi_h(x_i - u) \cdot \|x_i - u\|) \leq \lim_n c / (nh^d) < \infty, \tag{18}$$

where the first inequality holds for some constant $c \in \mathbb{R}^+$ because X is a compact set and because φ_h is bounded, non-negative, and has a finite support; and the second inequality holds because $nh^d = n^{4/(d+4)}$. The finiteness of the above summation means we can apply Kolmogorov's strong law of large numbers, which gives $\text{as-}\lim_n \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \|x_i - u\| - \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \|x_i - u\| \right) = 0$. Our next step is to compute this expectation. Note that

$$\begin{aligned}
\mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \|x_i - u\| \right) & \leq \int_{x \in \mathbb{R}^d} \varphi_h(x - u) \cdot \|x - u\| \cdot f_X(x) dx \\
& \leq h \int_{z \in \mathbb{R}^d} \varphi_1(z) \cdot z \cdot f_X(u + hz) dz \\
& \leq c \cdot h = c \cdot n^{-1/(d+4)}
\end{aligned} \tag{19}$$

where the second line makes the change of variables $z = (x - u)/h$, and the third line holds for some constant $c \in \mathbb{R}^+$ because the kernel has finite support and the density is continuous. The above expectation is non-negative, and so $\lim_n \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot \|x_i - u\| \right) = 0$. This proves the third result when combined with the outcome of Kolmogorov's strong law of large numbers. \square

Remark 7 It is important to note that the quantity $\gamma = \int_{\mathbb{R}^d} \varphi_1(z) dz$ is strictly greater than zero and finite because of the assumptions we have made.

4.3 Kernel Regression Estimator

We define a kernel regression estimate of S at the point u to be

$$\widehat{S}(u) = \left[\frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S_i \right] \cdot \left[\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \right]^{-1} \quad (20)$$

The following theorem proves the strong pointwise consistency of this estimator.

Theorem 5 *If $h = n^{-1/(d+4)}$, then $\text{as-lim}_n \widehat{S}(u) = S(u)$ for $u \in \text{int}(X)$.*

Proof Let $\kappa \in \mathbb{R}_+$ be the Lipschitz constant of S , and note that by Lipschitz continuity we have

$$\begin{aligned} \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S(u) \oplus w_i) &\subseteq \\ \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S_i \oplus \kappa \|x_i - u\| \mathbb{B}) &\subseteq \\ \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S(u) \oplus 2\kappa \|x_i - u\| \mathbb{B} \oplus w_i) &\quad (21) \end{aligned}$$

Corollary 2(c) and Lemma 2(a) give that $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S(u) = \gamma \cdot f_X(u) \cdot S(u)$, and Corollary 2(a) and Lemma 2(b) yield $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S(u) \oplus w_i) = \gamma \cdot f_X(u) \cdot S(u)$. Corollary 2(c) and Lemma 2(c) imply $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot \kappa \|x_i - u\| \mathbb{B} = 0$, and so Corollary 2(a) implies $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S(u) \oplus 2\kappa \|x_i - u\| \mathbb{B} \oplus w_i) = \gamma \cdot f_X(u) \cdot S(u)$. So applying the sandwich lemma to (21) yields $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot (S_i \oplus \kappa \|x_i - u\| \mathbb{B}) = \gamma \cdot f_X(u) \cdot S(u)$. Corollary 3 gives $\text{as-lim}_n \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S_i = \gamma \cdot f_X(u) \cdot S(u)$. Finally, using Corollary 2(d) and Lemma 2(a) imply that $\text{as-lim}_n \widehat{S}(u) = S(u)$. \square

Remark 8 The result above also applies to points $u \in \partial X$ when the set X is Clarke regular, and we briefly sketch the proof. The main difference is that the constant γ from Lemma 2 is replaced by a function $\zeta(u) = \int_{z \in T_X(u)} \varphi_1(z) dz$, where $T_X(u)$ is the tangent cone of X at the point u . (Observe that $\zeta(u) = \gamma$ for $u \in \text{int}(X)$ since $T_X(u) = \mathbb{R}^d$ for $u \in \text{int}(X)$.) The remaining analysis goes through after this change.

4.4 Algorithms to Compute Kernel Regression Estimator

The statistical consistency of our kernel regression estimator is a theoretical result, and numerical computation of this estimator needs some discussion. The key point is that the corresponding algorithm used to compute the estimator depends on the representation of the sets S_i that are in the measured data (u_i, S_i) for $i = 1, \dots, n$. Since the random sets S_i are RaTS, we only need to consider different representations of convex sets. Moreover, we focus our discussion on polytope representations since any compact convex set can be approximated arbitrarily well by polytopes [39].

If the sets S_i are each represented by polynomial time membership oracles, then

$$\frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S_i = \left\{ \frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot s_i : s_i \in S_i \text{ for } i = 1, \dots, n \right\}, \quad (22)$$

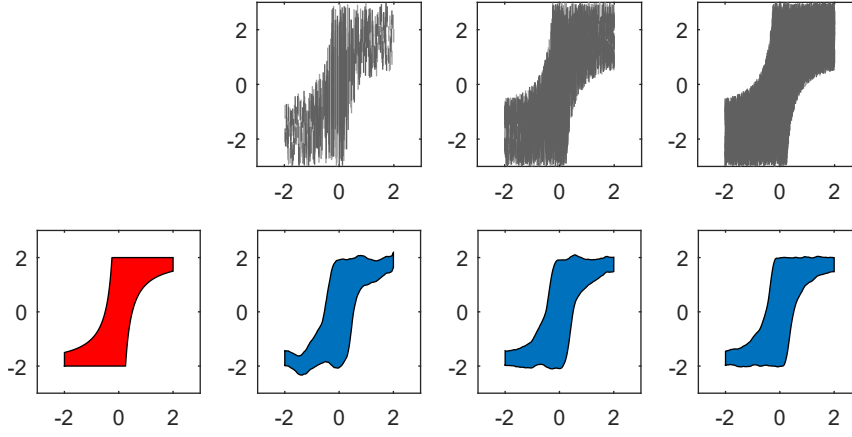


Fig. 1 The x -axes on the plots are u . The top row from left to right shows the noisy measurements (u_i, S_i) for $n = 10^2$, $n = 10^3$, and $n = 10^4$ data points, respectively, and the bottom row shows the set-valued function $S(u)$ being estimated and the corresponding estimates $\hat{S}(u)$ from our kernel regression estimator.

and so membership in the Minkowski sum can be determined in polynomial time. One example of when S_i is represented by a polynomial time membership oracle is if the measurements of S_i correspond to the polytopes $S_i = \{s_i : A_i s_i \leq b_i\}$, meaning the measurements are the values A_i, b_i in this example. Similar examples can be constructed for other classes of convex sets (e.g., semidefinite cone representable sets).

Zonotopes are polytopes defined as the Minkowski sum of line segments. Suppose the sets S_i are each represented by the zonotopes $S_i = \bigoplus_{k=1}^p w_{ik} \cdot I_k$, where I_k are line segments. Restated, the measurement are the w_{ik} and I_k in this setting. Then we have

$$\frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S_i = \bigoplus_{k=1}^p \left[\frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot w_{ik} \right] \cdot I_k, \quad (23)$$

and so the Minkowski sum is polynomial time computable for this representation.

Lastly, suppose the sets S_i are represented by the convex hull of a finite set of p_i vertices, meaning that $S_i = \text{co}(\{v_{i1}, \dots, v_{ip_i}\})$. In this setting the measurements are the vertices of each set S_i , and the Minkowski sum is given by

$$\frac{1}{n} \bigoplus_{i=1}^n \varphi_h(x_i - u) \cdot S_i = \text{co} \left(\left\{ \frac{1}{n} \sum_{i=1}^n \varphi_h(x_i - u) \cdot v_{ij_i} : \right. \right. \\ \left. \left. \text{for } j_i = 1, \dots, p_i \text{ and } i = 1, \dots, n \right\} \right). \quad (24)$$

This is a polynomial time computation since the number of vertices is finite.

4.5 Numerical Example

We conclude our discussion on kernel regression of set-valued functions with a numerical example to visually demonstrate the estimation problem being solved by our

estimator. Consider the set-valued function $S(u) : [-2, 2] \rightrightarrows \mathbb{R}$ that is given by

$$S(u) = \begin{cases} [-2, -\frac{2u+1}{u} + 2], & \text{if } u \in [-2, -\frac{1}{4}] \\ [-2, 2], & \text{if } u \in [-\frac{1}{4}, \frac{1}{4}] \\ [\frac{4u-1}{u} - 2, 2], & \text{if } u \in [\frac{1}{4}, 2] \end{cases} \quad (25)$$

This set-valued function is shown in the bottom-left of Fig. 1. Each measurement (u_i, S_i) for this example corresponds to a line segment S_i that is located at u_i , and the top row of Fig. 1 shows the measurements for $n = 10^2$, $n = 10^3$, and $n = 10^4$ data points, respectively. The bottom row shows the estimates $\hat{S}(u)$ of the set-valued function, and these estimates were computed using the measurements shown directly above the estimates. This example shows that as the number of measurements increases, the estimates $\hat{S}(u)$ converge pointwise to the actual set-valued function $S(u)$.

5 Inverse Approximate Optimization

Inverse optimization problems typically involve computing a set of parameters that make some measured solutions optimal [1, 8, 11, 15, 19, 23]. The inverse approximation optimization problem consists of a situation where noisy measurements of suboptimal solutions to an optimization problem are measured, and the goal is to estimate the amount of suboptimality of the solutions and estimate information about the objective and constraint functions defining the optimization problem generating the data. In principle, the VIA [11] and KKT [23] estimators can provide estimates of the desired quantities in this setting; but we show that their estimates are statistically inconsistent. As a result, we construct an estimator for inverse approximate optimization, prove its statistical consistency, and then discuss some possible generalizations.

5.1 Problem Setup

Consider a parametric convex optimization problem

$$V(u, \theta) = \min_x \{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\} \quad (26)$$

in which f, g are continuous functions that are convex in x for each fixed value of u and θ . We use the definition that ε -optimal solutions are those in the set

$$S(u, \varepsilon, \theta) = \varepsilon\text{-arg min}_x \{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\} = \{x : f(x, u, \theta) \leq V(u, \theta) + \varepsilon, g(x, u, \theta) \leq 0\}. \quad (27)$$

Our results are qualitatively the same if ε -optimal solutions are instead defined as $\varepsilon\text{-arg min}_x \{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\} = \{x : f(x, u, \theta) \leq V(u, \theta) + \varepsilon, g(x, u, \theta) \leq \varepsilon\}$. The difference is that the definition in (27) does not allow any constraint violation, while the alternative definition allows ε amount of constraint violation.

Now suppose that ε -optimal solutions of the optimization problem (26) generate the data $(u_i, y_i) \in U \times \mathbb{R}^p$ for $i = 1, \dots, n$, where: u_i are i.i.d. (vector-valued) random

variables distributed on the set $U \subseteq \mathbb{R}^d$; $y_i = x_i + w_i$, where x_i are i.i.d. (vector-valued) random variables distributed on $S(u_i, \varepsilon_0, \theta_0)$ with constants $\varepsilon_0 \in \mathbb{R}_+$ and $\theta_0 \in \mathbb{R}^p$; and w_i are i.i.d. (vector-valued) random variables with zero mean $\mathbb{E}(w) = 0$ and distributed on a known convex set W with finite support (which implies finite variance). We also assume the densities of w_i, x_i are strictly positive on the interior of their supports (i.e., $f_W(u) > 0$ for $u \in \text{int}(W)$ and $f_X(u|u_i) > 0$ for $u \in \text{int}(S(u_i, \varepsilon_0, \theta_0))$).

The inverse approximate optimization problem is to estimate $(\varepsilon_0, \theta_0)$ using the data (u_i, y_i) for $i = 1, \dots, n$. Note that we assume the functional forms of f, g are fixed. Let $E \subseteq \mathbb{R}_+$ and $\Theta \subseteq \mathbb{R}^p$ be known sets such that $\varepsilon_0 \in E$ and $\theta_0 \in \Theta$, and the intuition is that these sets represent prior knowledge that constrain the parameters and amount of solution suboptimality. The choice $E = \mathbb{R}_+$ and $\Theta = \mathbb{R}^p$ corresponds to a situation with no such prior knowledge. A so-called *identifiability condition* [13] is also needed. We assume that for all $(\varepsilon_0, \theta_0) \in E \times \Theta$ there exists $T(\varepsilon, \theta) \subseteq U$ such that $\mathbb{P}(u_i \in T(\varepsilon, \theta)) > 0$ and $S(u, \varepsilon, \theta) \not\subseteq S(u, \varepsilon_0, \theta_0)$ for $u \in T(\varepsilon, \theta)$ and $(\varepsilon, \theta) \in ([0, \varepsilon_0] \times \Theta) \setminus (\varepsilon_0, \theta_0)$. An identifiability condition (such as the one we have assumed) intuitively says that different parameters of the model produce different outputs.

5.2 Inconsistency of Existing Estimators

The VIA [11] and KKT [23] estimators have been shown [8] to be statistically inconsistent for the case when $\varepsilon_0 = 0$, but since these approaches minimize the amount of suboptimality of the measured data it is initially unclear without further analysis whether these approaches are inconsistent for problem instances with $\varepsilon_0 > 0$. The following result provides qualitative insights into the behavior of these estimators.

Proposition 4 *Let $r \in \mathbb{R}_+$ be a constant, and suppose $f = x^2$, $g = [x - 1; -x - 1]$, $\varepsilon_0 = 1$, $W = \{w : \|w\| \leq r\}$, and w_i, x_i are uniformly distributed. Then estimates $\hat{\varepsilon}$ generated by the VIA [11] and KKT [23] methods are such that $\liminf_n \hat{\varepsilon} > r/3$.*

Proof The KKT estimate is given by

$$\hat{\varepsilon} = \max \left\{ \frac{1}{n} \sum_{i=1}^n \langle g(y_i) \rangle_1^+, \frac{1}{n} \sum_{i=1}^n \langle g(y_i) \rangle_2^+, \frac{1}{n} \sum_{i=1}^n |2y_i + \langle \lambda_i \rangle_1 - \langle \lambda_i \rangle_2|, \right. \\ \left. \frac{1}{n} \sum_{i=1}^n |\langle \lambda_i \rangle_1 \cdot (y_i - 2)|, \frac{1}{n} \sum_{i=1}^n |\langle \lambda_i \rangle_2 \cdot (-y_i - 2)| \right\} \quad (28)$$

where these λ_i are the minimizers of the below optimization problem

$$\min \left\{ \frac{1}{n} \sum_{i=1}^n |2y_i + \langle \lambda_i \rangle_1 - \langle \lambda_i \rangle_2| + \langle \lambda_i \rangle_1 \cdot |x - 2| + \right. \\ \left. \langle \lambda_i \rangle_2 \cdot |-x - 2| \mid \lambda_i \geq 0, i = 1, \dots, n \right\}. \quad (29)$$

Since $\varepsilon\text{-arg min}_x \{f(x, u, \theta) \mid g(x, u, \theta) \leq 0\} = [-1, 1]$ under the hypothesis of this proposition, it holds that y_i are i.i.d. and have triangular distribution with lower limit $-r - 1$, upper limit $r + 1$, and mode 0. Hence the density of $c_i = \langle g(y_i) \rangle_1^+$ is given by

$$f_C(u) = \left(\frac{1}{2} + \frac{1}{r+1} - \frac{1}{2(r+1)^2} \right) \cdot \delta(u) + \left(\frac{r}{(r+1)^2} - \frac{1}{(r+1)^2} \cdot u \right) \cdot \mathbf{1}(u \in [0, r]), \quad (30)$$

where $\delta(u)$ is the Dirac delta function. So $\mathbb{E}(c_i) = \frac{r+1}{3}$, and the strong law of large numbers implies $\frac{r+1}{3} = \text{as-lim}_n \frac{1}{n} \sum_{i=1}^n c_i = \text{as-liminf}_n \frac{1}{n} \sum_{i=1}^n c_i \leq \text{as-liminf}_n \hat{\varepsilon}$.

The VIA estimate is given by $\hat{\varepsilon} = \frac{1}{n} \sum_{i=1}^n |\hat{\varepsilon}_i|$ where ε_i are the minimizers to

$$\min \left\{ \frac{1}{n} \sum_{i=1}^n |\hat{\varepsilon}_i| \mid 2y_i(x - y_i) \geq -\hat{\varepsilon}_i \text{ for } x \in [-1, 1], i = 1, \dots, n \right\} \quad (31)$$

However, observe that $2y_i(x - y_i) \geq -\hat{\varepsilon}_i$ for $x \in [-1, 1]$ simplifies to the constraint $-2(|y_i| + y_i^2) \geq -\hat{\varepsilon}_i$. Since the above optimization is minimizing each $|\hat{\varepsilon}_i|$, this means the constraint will be $\hat{\varepsilon}_i = 2(|y_i| + y_i^2)$ at optimality. Recall that as shown in the proof for KKT, the y_i have a triangular distribution with lower limit $-r - 1$, upper limit $r + 1$, and mode 0. This means $\mathbb{E}(|\hat{\varepsilon}_i|) = \frac{2r+2}{3} + \frac{(r+1)^2}{9}$. Applying the strong law of large numbers gives $\frac{2r+2}{3} + \frac{(r+1)^2}{9} = \text{as-lim}_n \frac{1}{n} \sum_{i=1}^n |\hat{\varepsilon}_i| = \text{as-lim}_n \hat{\varepsilon}$. \square

This proposition shows that existing approaches cannot distinguish between noise in measurements versus suboptimality of the solutions. The reason is that these approaches are minimizing an incorrect error metric: They minimize the amount of suboptimality of the measured data, and this is an incorrect error metric when the measured data is noisy because the noise increases the suboptimality of the measured data. Moreover, this indistinguishability of existing approaches is unbounded in the sense that as the noise variance increases then their estimates of suboptimality increase in an unbounded way. Such behavior is undesirable, and in fact the above result gives the following corollary on the statistical properties of VIA and KKT.

Corollary 4 *The VIA [11] and KKT [23] estimators are statistically inconsistent.*

Proof By definition an estimator is consistent for a class of models if and only if it is consistent for each model in that class. Thus to show inconsistency of VIA and KKT it suffices to show inconsistency for a single model. The above proposition establishes inconsistency of VIA and KKT for a particular model because $\varepsilon_0 = 1$ while $\text{as-liminf}_n \hat{\varepsilon} > r/3$, meaning these approaches are inconsistent when $r > 3$. \square

5.3 Approximate Bilevel Programming (ABP) Estimator

To correct the indistinguishability (between suboptimality of solutions and noise in measurements) problem faced by existing approaches, we instead propose an estimator that explicitly models the measured data as consisting of a suboptimal solution added to noise. More specifically, we propose the following statistical estimator

$$(\check{\varepsilon}, \check{\theta}) \in \arg \min \left\{ \frac{1}{n} \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) + \lambda \cdot \varepsilon \mid \varepsilon \in E, \theta \in \Theta \right\} \quad (32)$$

where $\lambda \in \mathbb{R}_+$ and ρ is the squared distance function defined in the preliminaries. It is also useful to consider estimators defined as approximate solutions to the above optimization problem. Let $z \in \mathbb{R}_+$ be a nonnegative value, and define the estimates

$$(\hat{\varepsilon}, \hat{\theta}) \in \left\{ \varepsilon \in E, \theta \in \Theta : \frac{1}{n} \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) + \lambda \cdot \varepsilon \leq \frac{1}{n} \sum_{i=1}^n \rho(y_i, S(u_i, \check{\varepsilon}, \check{\theta}) \oplus W) + \lambda \cdot \check{\varepsilon} + z \right\}. \quad (33)$$

For notational convenience, we will call this estimator the ABP estimator. Note these estimates are defined as being any z -argmin of the optimization problem (32).

Theorem 6 *The ABP estimator is strongly statistically consistent, meaning we have $\text{as-lim}_n(\hat{\varepsilon}, \hat{\theta}) = (\varepsilon_0, \theta_0)$ whenever $\lambda = 1/n$ and $\lim_n(n \cdot z) = 0$.*

Proof Our first step is to show $\rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)$ satisfies certain continuity properties. Note $\{x : g(x, u, \theta) \leq 0\}$ is continuous by Example 5.10 of [35], and so $V(u, \theta)$ is continuous by the Berge maximum theorem [10]. Noting $\rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) = \min\{\|y_i - \hat{y}_i\|^2 \mid \hat{y}_i = \hat{x}_i + \hat{\varepsilon}_i, \hat{\varepsilon}_i \in W, f(\hat{x}_i, u_i, \theta) \leq V(u_i, \theta) + \varepsilon, g(\hat{x}_i, u_i, \theta) \leq \varepsilon\}$, we can apply the Berge maximum theorem [10] since this feasible set is osc by Example 5.8 of [35]: This implies $\rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)$ is lower semicontinuous on $E \times \Theta$.

Next note that the estimate (ε, θ) also minimizes the optimization problem

$$\min \left\{ \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) + n\lambda \cdot \varepsilon \mid \varepsilon \in E, \theta \in \Theta \right\} \quad (34)$$

But $n\lambda = 1$ by assumption, and so the objective of (34) is nondecreasing. Hence $\mathbb{P}(\text{e-lim}_n \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) = \sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)) = 1$ by Proposition 7.4 of [35], where e-lim is the epi-limit [35]. We use contradiction to prove the claim $\mathbb{P}(\sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) = \delta((\varepsilon, \theta), \{(\varepsilon_0, \theta_0)\})$ for $(\varepsilon, \theta) \in [0, \varepsilon_0] \times \Theta) = 1$ where δ is the indicator function defined in the preliminaries. If this claim is false, then there exists $(\underline{\varepsilon}, \underline{\theta}) \in [0, \varepsilon_0] \times \Theta$ such that $\mathbb{P}(\sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \underline{\varepsilon}, \underline{\theta}) \oplus W) < +\infty) > 0$. Note $(\underline{\varepsilon}, \underline{\theta}) \neq (\varepsilon_0, \theta_0)$ because otherwise a strict inequality could not hold since all the relevant functions are nonnegative. But $\mathbb{P}(\sum_{i=1}^n \rho(y_i, S(u_i, \underline{\varepsilon}, \underline{\theta}) \oplus W) = 0) = (1 - \mathbb{P}(u \in T(\underline{\varepsilon}, \underline{\theta})))^n$ and so $\sum_{n=1}^\infty \mathbb{P}(\sum_{i=1}^n \rho(y_i, S(u_i, \underline{\varepsilon}, \underline{\theta}) \oplus W) = 0) < +\infty$ since $\mathbb{P}(u \in T(\underline{\varepsilon}, \underline{\theta})) > 0$ by assumption. Using Borel-Cantelli lemma implies that $\mathbb{P}(\sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \underline{\varepsilon}, \underline{\theta}) \oplus W) < +\infty) = 0$, which is a contradiction. This shows the desired claim, and so we conclude $\sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)$ epi-converges almost surely to $\delta((\varepsilon, \theta), \{(\varepsilon_0, \theta_0)\})$ on the domain $[0, \varepsilon_0] \times \Theta$.

Next consider the optimization problem

$$\min \left\{ \sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) + \varepsilon \mid \varepsilon \in E, \theta \in \Theta \right\} \quad (35)$$

Note $(\varepsilon_0, \theta_0)$ is feasible for both (34) and (35), and so the minimums of (34) and (35) are both less than or equal to ε_0 . This means $\varepsilon > \varepsilon_0$ cannot minimize (35). Furthermore, we showed above that $\sup_n \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W) = \delta((\varepsilon, \theta), \{(\varepsilon_0, \theta_0)\})$ on $[0, \varepsilon_0] \times \Theta$ almost surely. Hence the (unique) minimizer of (35) is $(\varepsilon_0, \theta_0)$, and the minimum value of (35) is ε_0 . But the optimization problem (34) epi-converges almost surely to (35) on $[0, \varepsilon_0] \times \Theta$ since: $\sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)$ epi-converges almost surely to $\delta((\varepsilon, \theta), \{(\varepsilon_0, \theta_0)\})$ on $[0, \varepsilon_0] \times \Theta$ as shown above, and the other terms of the optimization problem are a constant linear term in the object ε or fixed constraints $\varepsilon \in E, \theta \in \Theta$. The result now follows from Theorem 7.33 of [35]. \square

The above result concerns almost sure convergence of the ABP estimates $(\hat{\varepsilon}, \hat{\theta})$ to the actual parameters $(\varepsilon_0, \theta_0)$, but a related question is whether the corresponding solution set estimates $S(u, \hat{\varepsilon}, \hat{\theta})$ converge to the actual solution sets $S(u, \varepsilon_0, \theta_0)$. Our semicontinuous mapping theorem can be used to establish almost sure convergence of the solution set estimates, and this argument leads to the the following corollary.

Corollary 5 *We have that $\text{as-lim sup}_n S(u, \hat{\varepsilon}, \hat{\theta}) \subseteq S(u, \varepsilon_0, \theta_0)$ for $u \in U$. If $\varepsilon_0 > 0$ or $f(\cdot, u, \theta)$ is strictly convex in x , then $\text{as-lim}_n S(u, \hat{\varepsilon}, \hat{\theta}) = S(u, \varepsilon_0, \theta_0)$ for $u \in U$.*

Proof The above proof established that $S(u, \varepsilon, \theta)$ is osc in ε, θ . And so the first part of the corollary follows by the semicontinuous mapping theorem. If $\varepsilon_0 > 0$ then $S(u, \varepsilon_0, \theta_0)$ is continuous at $(\varepsilon_0, \theta_0)$ by Example 5.10 of [35]. If $f(\cdot, u, \theta)$ is strictly convex in x , then $S(u, 0, \theta_0)$ is single-valued [35]. Hence $S(u, 0, \theta_0)$ is continuous because a single-valued, osc, and locally bounded function is continuous [35]. Thus the second part of the corollary follows from the semicontinuous mapping theorem. \square

5.4 Algorithms to Compute ABP Estimator

The above results concern the statistical properties of the ABP estimator, and in this subsection we discuss numerical computation of the estimator. The ABP estimator can be seen to be an approximate (in the sense that the solution sets have ε possibly greater than zero) bilevel program, and so the relevant starting point is numerical methods for bilevel programs. Bilevel programs are optimization problems in which some decision variables are solutions to optimization problems, and so the lower level problem refers to this optimization problem that constrains the decision variables. There are three categories of approaches. The first is to replace the lower level problem with their KKT conditions [2, 16], and for certain problem structures this resulting problem can be rewritten as mixed-integer linear or quadratic programs that can be numerically solved very quickly [6]. The second is to upper bound the objective function of the lower level problem by its value function [32, 45].

Here we describe how a third approach that upper bounds the objective function of the lower level problem by its dual function [7, 8] can be used to compute the ABP estimator. If $h(u_i, \theta, \lambda)$ is the Lagrangian dual function corresponding to (26), then under mild conditions ensuring zero duality gap the ABP estimator is given by

$$\begin{aligned} (\check{\varepsilon}, \check{\theta}) \in \arg \min & \frac{1}{n} \sum_{i=1}^n \|y_i - \hat{x}_i\|^2 + \lambda \cdot \varepsilon \\ \text{s.t. } & f(\hat{x}_i, u_i, \theta) \leq h(u_i, \theta, \lambda_i) + \varepsilon \\ & g(\hat{x}_i, u_i, \theta) \leq 0 \\ & \lambda_i \geq 0, \varepsilon \in E, \theta \in \Theta \end{aligned} \tag{36}$$

This duality-based reformulation can be numerically solved by two different algorithms [7, 8], which we briefly describe here. More details can be found in the corresponding references, and both algorithms assume the sets E, Θ are compact.

Since the reformulation (36) is a convex optimization problem for fixed (ε, θ) , one algorithm [8] for computing ABP is to: discretize the set $E \times \Theta$ into a finite set $\Delta = \{(\varepsilon_1, \theta_1), \dots, (\varepsilon_k, \theta_k)\}$ such that it forms a set covering with balls of a prescribed radius, compute the minimum objective function value of (36) for $(\varepsilon, \theta) \in \Delta$ (which we call $Q(\varepsilon, \theta)$), and then choose estimates $(\hat{\varepsilon}, \hat{\theta}) = \arg \min \{Q(\varepsilon, \theta) \mid (\varepsilon, \theta) \in \Delta\}$. A result from [8] implies that estimates chosen using this *enumeration algorithm* satisfy the assumptions of Theorem 6, which is sufficient for statistical consistency.

A second algorithm [7] replaces the Lagrangian dual by a numerically computed dual. Partial dualization is used to define a regularized dual function (RDF)

$$h_\mu(u, \theta, \lambda) = \min_x \{ \mu \cdot \|x\|^2 + f(x, u, \theta) + \lambda^\top g(x, u, \theta) \mid x \in X \}. \quad (37)$$

Here, X is any compact set defined such that $\{x : \exists(u, \theta) \in U \times \Theta \text{ s.t. } g(x, u, \theta) \leq 0\} \subseteq \text{int}(X)$. The intuition is that X is a set that contains all the feasible sets of (26) within its interior. When g does not depend on (u, θ) , we can choose $X = \{x : l_i - 1 \leq x_i \leq u_i + 1\}$ with $u_i = \max\{x_i \mid g(x) \leq 0\}$ and $l_i = \min\{x_i \mid g(x) \leq 0\}$ that are computed by solving convex optimization problems. Many applications of inverse approximate optimization consist of such a setting where the feasible set is independent of the inputs u or the parameters θ . The benefit of the RDF is it can be numerically computed because it is a convex optimization problem, and that its gradient

$$\begin{aligned} \nabla_\theta h_\mu(u, \theta, \lambda) &= \nabla_\theta f(x, u, \theta) + \lambda^\top \cdot \nabla_\theta g(x, u, \theta) \\ \nabla_\lambda h_\mu(u, \theta, \lambda) &= g(x, u, \theta) \end{aligned} \quad (38)$$

always exists when $\mu > 0$. In contrast, the Lagrangian dual is usually only directionally differentiable but not differentiable. The algorithm proceeds by using a nonlinear numerical solver to solve a sequence of optimization problems in which μ goes to 0.

5.5 Numerical Example

We next consider a numerical example to visually compare estimates of the approximate solution set $S(u_i, \varepsilon_0, \theta_0)$ produced by our ABP estimator to those produced by the VIA [11] and KKT [23] estimators. Suppose $x \in \mathbb{R}$, and consider the setting where $f = -(\theta + u) \cdot x$, $g = [x - 2; -x - 2]$, $\varepsilon_0 = 1$, $W = \{w : \|w\| \leq 1\}$, u is uniformly distributed on $U = \{u : \|u\| \leq 2\}$, x_i is uniformly distributed on $S(u_i, \varepsilon_0, \theta_0)$, $E = \{\varepsilon : 0.1 \leq \varepsilon \leq 10\}$, and $\Theta = \{\theta : -2 \leq \theta \leq 2\}$. The solution set $S(u, \varepsilon_0, \theta_0)$ in this setting is shown in the left column of Fig. 2. Each measurement (u_i, y_i) for this example is a point, and the top row of Fig. 2 shows the measurements for $n = 10^1$, $n = 10^2$, and $n = 10^3$ data points, respectively. The rows below show the estimated (using the measurements shown above) solution set as computed by ABP, KKT, and VIA, respectively. This example shows that as the number of measurements increases, the solution set estimated by ABP (KKT and VIA) converges (do not converge) to the actual solution set. This statistical behavior is expected given our theoretical results on the strong consistency of ABP and the statistical inconsistency of KKT and VIA.

5.6 Related Inverse Optimization Problems

In our problem setup, the measurement noise w_i had a distribution with a finite support. However, noise models commonly used in statistics include distributions with unbounded support but finite variance. The canonical example is w_i that are jointly Gaussian with zero mean and finite covariance. A heuristic approach for distributions with unbounded support is to use our ABP estimator with the choices of

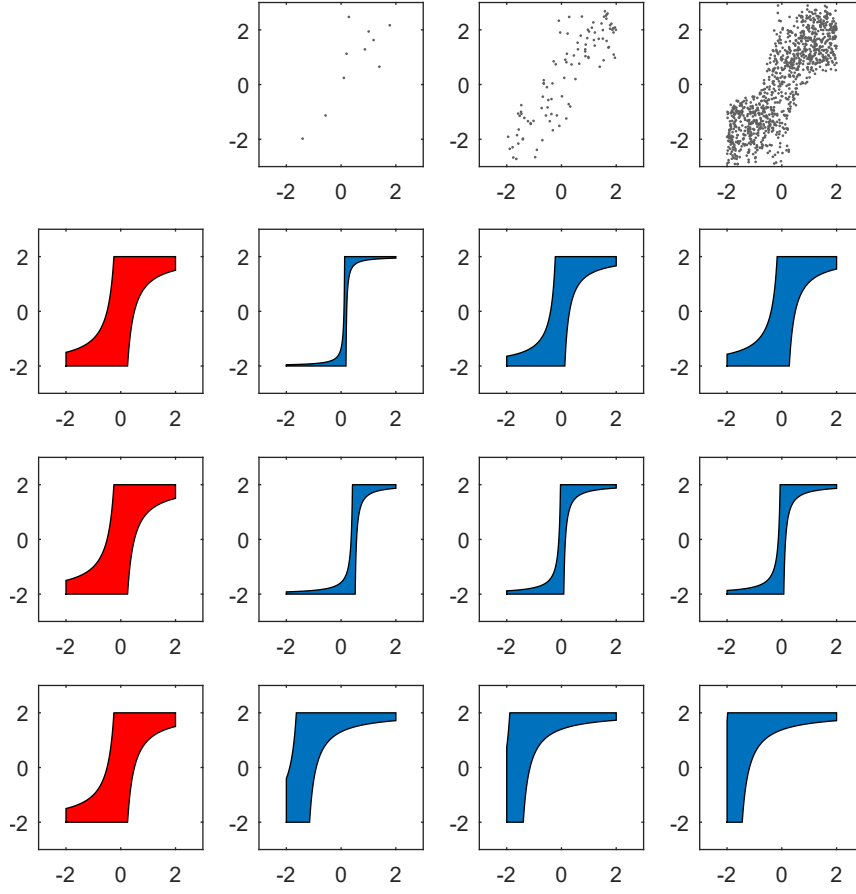


Fig. 2 The x -axes on the plots are u . For the solution set $S(u, \varepsilon, \theta) = \varepsilon \cdot \arg \min_x \{-(\theta + u) \cdot x \mid -2 \leq x \leq 2\}$, the left column shows the actual $S(u, \varepsilon_0, \theta_0)$ when $\varepsilon_0 = 1$ and $\theta_0 = 0$. The top row from left to right shows the noisy measurements (u_i, y_i) for $n = 10^1$, $n = 10^2$, and $n = 10^3$ data points, respectively, and the rows below show the estimated solution set as computed by ABP, KKT, and VIA, respectively.

$W = (2 \log n)^{1/2} \cdot \Sigma$ for sub-Gaussian distributions (i.e., distributions bounded from above by a jointly Gaussian random variable) and $W = ((2 \log n)^{1/2} + \log n) \cdot \Sigma$ for sub-exponential distributions (i.e., distributions with exponentially decaying tails), where $\Sigma = \mathbb{E}(w w^T)$ is the covariance matrix of w_i . The reason for this suggested heuristic is these choices of W are analogous to bounds on the maximum expected values of sub-Gaussian and sub-exponential random variables [14].

Since the ABP estimator is a heuristic in this setting, an obvious topic is to design a statistically consistent estimator for inverse approximate optimization problems with unbounded noise. Maximum likelihood estimation is arguably the most natural approach because otherwise it is difficult to distinguish between noise and suboptimality of solutions. More specifically, suppose we consider the original problem setup but with the changes that x_i are uniformly distributed within $S(u_i, \varepsilon_0, \theta_0)$,

and that the w_i are distributed according to some known density $f_W(u)$. Then the maximum likelihood estimator (MLE) for this modified problem setup is given by

$$(\varepsilon_{\text{mle}}, \theta_{\text{mle}}) \in \arg \min \left\{ -\frac{1}{n} \sum_{i=1}^n \log \int_{x \in S(u_i, \varepsilon, \theta)} f_W(y_i - x) dx + \frac{1}{n} \sum_{i=1}^n \log \int_{x \in S(u_i, \varepsilon, \theta)} dx \mid \varepsilon \in E, \theta \in \Theta \right\}. \quad (39)$$

This optimization problem has a challenging structure in which the domains of integration depend upon the decision variables, and presents an opportunity for the further study of designing numerical algorithms to solve such optimization problems. We do note that for fixed (ε, θ) , the integrals in the objective can be numerically computed in polynomial time using hit-and-run techniques for sampling from convex sets [28, 41]. And so the enumeration algorithm we described earlier for the ABP estimator could be easily modified to solve this MLE problem.

Remark 9 The ABP and MLE estimators are actually qualitatively the same. The $\frac{1}{n} \sum_{i=1}^n \rho(y_i, S(u_i, \varepsilon, \theta) \oplus W)$ term in ABP and the $-\frac{1}{n} \sum_{i=1}^n \log \int_{x \in S(u_i, \varepsilon, \theta)} f_W(y_i - x) dx$ term in MLE both penalize estimates in which the measured solutions y_i are far from the solution sets $S(u_i, \varepsilon, \theta)$, and the $\lambda \cdot \varepsilon$ term in ABP and the $\frac{1}{n} \sum_{i=1}^n \log \int_{x \in S(u_i, \varepsilon, \theta)} dx$ term in MLE both penalize estimates that generate large solution sets.

In the two inverse approximate optimization problem setups considered above, we assumed the approximate solutions x_i were drawn from the solution sets $S(u_i, \varepsilon_0, \theta_0)$ according to some distribution. However, another modified problem setup would be to assume the x_i were chosen from the solution sets by solution of another optimization problem. This kind of setup corresponds to a scenario in which the x_i are solutions to an optimistic bilevel optimization problem with unique solutions:

$$x_i = \min \{F(u_i, x_i, z_i) \mid x_i \in S(u_i, \varepsilon_0, \theta_0), G(u_i, x_i, z_i) \leq 0\}. \quad (40)$$

In this case, the estimation procedure can be posed as a least squares problem

$$(\varepsilon_{\text{ble}}, \theta_{\text{ble}}) \in \arg \min \frac{1}{n} \sum_{i=1}^n \|y_i - x_i\|^2$$

$$\text{s.t. } x_i = \min \{F(u_i, x_i, z_i) \mid x_i \in S(u_i, \varepsilon, \theta), G(u_i, x_i, z_i) \leq 0\} \quad (41)$$

$$\varepsilon \in E, \theta \in \Theta.$$

This is a challenging multi-level optimization problem and presents an opportunity for the further study of designing numerical algorithms to solve such optimization problems. We do note that for fixed (ε, θ) , this becomes a convex optimization problem. And so the enumeration algorithm we described earlier for the ABP estimator could be easily modified to solve this least squares problem.

6 Conclusion

In this paper, we used variational analysis to develop a framework for statistics with set-valued functions, and then applied this framework to two estimation problems. We constructed and studied a kernel regression estimator for set-valued functions and

an estimator for the inverse approximate optimization problem. The area of statistics with set-valued functions remains largely unexplored with many remaining problems. One question is the design of numerical representations of sets and set-valued functions. Though constraint representations of sets are pervasive, numerical machinery like epi-splines [37] may offer greater representational flexibility. Another question is the development of numerical algorithms to solve optimization problems that arise in statistical estimation for set-valued functions. Related inverse optimization problems lead to formulations (39) and (41) with structures that are not well-studied from the perspective of numerical optimization. Further study of statistics with set-valued functions will require developing new numerical methods and optimization theory.

References

1. Ahuja, R., Orlin, J.: Inverse optimization. *Operations Research* **49**(5), 771–783 (2001)
2. Allende, G., Still, G.: Solving bilevel programs with the KKT-approach. *Mathematical programming* **138**(1-2), 309 (2013)
3. Artstein, Z., Vitale, R.: A strong law of large numbers for random compact sets. *The Annals of Probability* pp. 879–882 (1975)
4. Aswani, A., Bickel, P., Tomlin, C.: Regression on manifolds: Estimation of the exterior derivative. *The Annals of Statistics* pp. 48–81 (2011)
5. Aswani, A., Gonzalez, H., Sastry, S., Tomlin, C.: Provably safe and robust learning-based model predictive control. *Automatica* **49**(5), 1216–1226 (2013)
6. Aswani, A., Kaminsky, P., Mintz, Y., Flowers, E., Fukuoka, Y.: Behavioral modeling in weight loss interventions. Available at SSRN: <https://ssrn.com/abstract=2838443> (2016). Submitted
7. Aswani, A., Ouattara, A.: Duality approach to bilevel programs with a convex lower level. *arXiv preprint arXiv:1608.03260* (2016)
8. Aswani, A., Shen, Z.J., Siddiq, A.: Inverse optimization with noisy data. *arXiv preprint arXiv:1507.03266* (2015)
9. Aumann, R.J.: Integrals of set-valued functions. *Journal of Mathematical Analysis and Applications* **12**(1), 1–12 (1965)
10. Berge, C.: *Topological Spaces: Including a Treatment of Multi-valued Functions, Vector Spaces, and Convexity*. Courier Corporation (1963)
11. Bertsimas, D., Gupta, V., Paschalidis, I.: Data-driven estimation in equilibrium using inverse optimization. *Mathematical Programming* **153**(2), 595–633 (2015)
12. Bickel, P.: On adaptive estimation. *The Annals of Statistics* pp. 647–671 (1982)
13. Bickel, P., Doksum, K.: *Mathematical Statistics: Basic Ideas And Selected Topics*, vol. 1, 2nd edn. Pearson Prentice Hall (2006)
14. Boucheron, S., Lugosi, G., Massart, P.: *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford (2013)
15. Chan, T., Craig, T., Lee, T., Sharpe, M.: Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research* (2014)
16. Dempe, S., Dutta, J.: Is bilevel programming a special case of a mathematical program with complementarity constraints? *Mathematical programming* **131**(1-2), 37–48 (2012)
17. Devroye, L., Wise, G.: Detection of abnormal behavior via nonparametric estimation of the support. *SIAM Journal on Applied Mathematics* **38**(3), 480–488 (1980)
18. Dupacová, J., Wets, R.: Asymptotic behavior of statistical estimators and of optimal solutions of stochastic optimization problems. *The annals of statistics* pp. 1517–1549 (1988)
19. Esfahani, P.M., Shafieezadeh-Abadeh, S., Hanasusanto, G.A., Kuhn, D.: Data-driven inverse optimization with incomplete information. *arXiv preprint arXiv:1512.05489* (2015)
20. Geffroy, J.: Sur un probleme destination géométrique. *Publ. Inst. Statist. Univ. Paris* **13**, 191–210 (1964)
21. Geyer, C.J.: On the asymptotics of constrained M-estimation. *The Annals of Statistics* pp. 1993–2010 (1994)

22. Guntuboyina, A.: Optimal rates of convergence for convex set estimation from support functions. *The Annals of Statistics* **40**(1), 385–411 (2012)
23. Keshavarz, A., Wang, Y., Boyd, S.: Imputing a convex objective function. In: *IEEE Multi-Conference on Systems and Control*, pp. 613–619 (2011)
24. Knight, K., Fu, W.: Asymptotics for lasso-type estimators. *The Annals of Statistics* pp. 1356–1378 (2000)
25. Korostelev, A., Simar, L., Tsybakov, A.: Efficient estimation of monotone boundaries. *The Annals of Statistics* pp. 476–489 (1995)
26. Kudo, H.: Dependent experiments and sufficient statistics. *Natural Science Report. Ochanomizu University* **4**, 905–927 (1953)
27. Lachout, P., Liebscher, E., Vogel, S.: Strong convergence of estimators as ε -minimisers of optimisation problems of optimisation problems. *Annals of the Institute of Statistical Mathematics* **57**(2), 291–313 (2005)
28. Lovász, L., Vempala, S.: Hit-and-run from a corner. *SIAM Journal on Computing* **35**(4), 985–1005 (2006)
29. Matheron, G.: Random sets and integral geometry. John Wiley & Sons (1975)
30. Molchanov, I.: Theory of random sets. Springer Science & Business Media (2006)
31. Noda, K.: Estimation of a regression function by the parzen kernel-type density estimators. *Annals of the Institute of Statistical Mathematics* **28**(1), 221–234 (1976)
32. Outrata, J.: On the numerical solution of a class of Stackelberg problems. *Zeitschrift für Operations Research* **34**(4), 255–277 (1990)
33. Patschkowski, T., Rohde, A., et al.: Adaptation to lowest density regions with application to support recovery. *The Annals of Statistics* **44**(1), 255–287 (2016)
34. Rényi, A., Sulanke, R.: Über die konvexe hülle von n zufällig gewählten punkten. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **2**(1), 75–84 (1963)
35. Rockafellar, R., Wets, R.: *Variational Analysis*, 3rd edn. Springer (2009)
36. Royset, J.: *Approximations and solution estimates in optimization* (2016)
37. Royset, J., Wets, R.: Multivariate epi-splines and evolving function identification problems. *Set-Valued and Variational Analysis* **24**(4), 517–545 (2016)
38. Salinetti, G., Wets, R.: On the convergence in distribution of measurable multifunctions (random sets) normal integrands, stochastic processes and stochastic infima. *Mathematics of Operations Research* **11**(3), 385–419 (1986)
39. Schneider, R.: *Convex bodies: the Brunn–Minkowski theory*. 151. Cambridge University Press (1993)
40. Schölkopf, B., Platt, J., Shawe-Taylor, J., Smola, A., Williamson, R.: Estimating the support of a high-dimensional distribution. *Neural computation* **13**(7), 1443–1471 (2001)
41. Smith, R.: Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research* **32**(6), 1296–1308 (1984)
42. van der Vaart, A.: *Asymptotic Statistics*. Cambridge University Press (2000)
43. Wand, M., Jones, M.: *Kernel Smoothing*. Taylor & Francis (1994)
44. Weil, W.: An application of the central limit theorem for Banach-space-valued random variables to the theory of random sets. *Probability Theory and Related Fields* **60**(2), 203–208 (1982)
45. Ye, J., Zhu, D.: Optimality conditions for bilevel programming problems. *Optimization* **33**(1), 9–27 (1995)